arXiv:1611.00502v1 [math.CO] 2 Nov 2016

A Simple Algorithmic Proof of the Symmetric Lopsided Lovász Local Lemma

Lefteris Kirousis
John Livieratos
Department of Mathematics
National and Kapodistrian University of Athens
emails: lkirousis@math.uoa.gr, jlivier89@math.uoa.gr

November 3, 2016

Abstract

We provide a simple algorithmic proof for the symmetric *Lopsided* Lovász Local Lemma, a variant of the classic Lovász Local Lemma, where, roughly, only the degree of the negatively correlated undesirable events counts. Our analysis refers to the algorithm by Moser in 2009, however it is based on a simple application of the probabilistic method, rather than a counting argument, as are most of the analyses of algorithms for variants of the Lovász Local Lemma.

Introduction

The Lovász Local Lemma (LLL) first appeared in 1975 in a paper by Erdős and Lovász [1]. In its original form, LLL provides a necessary condition for the possibility to avoid a number of undesirable events given that there is a small constant upper bound on the product of the maximum probability of the events with the maximum number of other events any given event depends on. This form is known as *symmetric* LLL.

There is also a more general, *non-symmetric*, form where the necessary condition asks for the existence of a family of numbers in $(0, 1)$, one for each event, such that for each event E its probability is bounded from above by a specified expression of the numbers that correspond to the neighbors of E. See e.g., the exposition paper by Szegedy [2] for the necessary background.

The proofs of these versions were non-constructive. The first algorithmic proof, that entailed an unexpectedly simple algorithm, was given much later by Moser [3] (work on constructive approaches was previously done by Alon [4], Beck [5], Srinivasan [6] and others). The initial proof of Moser was only for the application of LLL in the satisfiability problem. It was based on a counting argument that became known as the "entropic method": one bounds the entropy (cardinality) of structures that witness that the algorithm lasts for too many steps (see [7]). For an elegant presentation of this result based on a direct probabilistic proof, see Spencer [8]. Closely related to Spencer's proof is the proof in Giotis et al. [9].

The so called *lopsided* version of the Lovász Local Lemma (LLLL), where dependency of undesirable events is restricted to negative correlation, was proved, in a non-constructive way, by Erdős and Spencer in 1991 [10]. This version as well has a more general non-symmetric variant and has been used in, among others, the satisfiability problem: instead of declaring two clauses of a formula dependent on each other when they share a variable, one can consider clauses that are in conflict with each other (i.e. have opposing literals of the same variable). Berman et al. [11] and Gebauer et al. [12], [13] have successfully used it to bound the number of occurrences per variable that can be allowed in a formula, while guaranteeing the existence of a satisfying assignment.

An algorithmic proof for the non-symmetric form, of both non-lopsided and lopsided versions, was given by Moser and Tardos [14]. The proof assumed that the events are expressed in terms of a family of independent variables. Specifically, it is assumed that each event depends on a subset of the independent variables that comprises the *scope* of the event; the dependency graph is given in terms of how the scopes of the events are related. This proof as well is based on the "entropic method".

Kolipaka and Szegedy [15] describe an algorithmic approach where the events are not assumed to be expressed in terms of independent variables. In their framework, they prove the most general form of LLL, the Shearer version [16], which provides a necessary and sufficient condition for LLL.

More recently, Harvey and Vondrák [17] provided an algorithmic proof of the general non-symmetric form of lopsided LLL, also without assuming that the events are expressed in terms of variables. Their proof was through the Shearer version.

It is worth pointing out that very recently Harris [18] claimed that the lopsided version of LLL in the variable framework can be stronger than the general Shearer version, if a suitable dependency graph is defined in the variable framework (this variable version was analyzed by Harris in [18], by the entropic method). The crucial point in their proof is that the dependency graph they define is essentially directed; directed dependency graphs cannot be defined in the framework without variables.

In this work, we deal with the lopsided LLL in the variable framework. We introduce a simple notion of *directed* dependency graph inspired by the undirected notion of lopsidependency of Gebauer et al. [13]. Our proof is based on the probabilistic method and not on a counting argument. Also it does not necessitate proving first the complex (but general) version of Shearer. However, we restricted our analysis to the symmetric case, where the product of the maximum probability over all events times the outgoing degree of the directed dependency graph (incremented by one) is assumed to be bounded by a constant. Although this version is less general than the non-symmetric versions, the lack of generality is compensated, we believe, by the simplicity of the analysis, which reflects the elegance of the algorithm of Moser [3] that we use. We consider this simplicity in comparison to some of the extant work as the primary contribution of this work; it is an outcome of using the probabilistic method (Spencer's first Durango lecture [19] vividly describes the power of this method). We also consider interesting the introduction of a simple variable-dependent directed notion of lopsidependency in Section 2; a notion that constitutes a further indication that for some applications, the variable framework might be advantageous over the no-variables framework of general probability spaces. Achlioptas and Il-

iopoulos [20] present a completely abstract framework for LLL, interestingly not in terms of probability spaces, that covers also the lopsided version with a directed dependency graph; for the proof they use counting arguments (entropic method).

Let us also point out that the difference of the algorithm by Moser [3], which we use, from the subsequent one of Moser and Tardos [14] is that the former, at each phase, resamples an occurring event that is dependent on the previously examined one, in contrast with the latter, which at each phase resamples an arbitrary event not necessarily depending on the previously examined one.

In Section 1, we formally present the variable framework. In Section 2, we formally present the various lopsidependency notions and show how they are related. Finally in Section 3, we present our main result.

1 The variable framework

We will work in what is known as the *variable framework*, which was first used in [3], [14]. Let X_i, $i = 1, ..., l$ be mutually independent random variables on a common probability space, taking values in the finite sets D_i, $i = 1, ..., l$, respectively.

An evaluation of the random variables is an l-ary vector $\alpha = (a_1, ..., a_l)$, where $a_i \in D_i$, for all $i \in \{1, ..., l\}$. Let $\Omega = \prod_{i=1}^{l} D_i$ be the probability space of all evaluations for the variables.

Let E_j, $j = 1, ..., m$, be a sequence of events. We define the *scope* e^j of event E_j to be the minimal subset of variables such that one can determine whether E_j is satisfied or not knowing only their values, i.e., event E_j depends only on the values of the variables of e^j. Events are assumed to be ordered according to their index.

The events E_j are considered "undesirable", i.e., the objective is to design a randomized algorithm that will return an evaluation α of the variables X_i for which none of the events E_j holds.

2 Lopsidependent events

We will first give a notion of dependency between a pair of events, that we call *Variable-dependent Directed Lopsidependency*, VDL in short (depending on the context, VDL may also stand for "Variable-dependent Directed Lopsidependent" –an adjective rather than a noun).

Definition 1. *Let E_i, E_j be events, $i, j \in \{1, ..., m\}$. We say that E_j is VDL on E_i if:*

1. *there exists an evaluation α that makes E_i and $\overline{E_j}$ occur and*
2. *there exists an evaluation β that differs from α only in variables in e^i, that makes E_j occur.*

Intuitively, E_j is VDL on E_i if the effort to undo the undesirable E_i results in E_j happening, although it did not happen before.

Notice that an event E_i can never be VDL on itself, nor on any event whose scope shares no variables with e^i.

The binary relation VDL defines a simple *directed* graph $G = (V, E)$, the *VDL graph* of events $E_1, ..., E_m$, where $V = \{1, ..., m\}$ and

$$E = \{(i, j) \mid E_j \text{ is VDL on } E_i\}.$$

For $i = 1, ..., m$, let $\Gamma(E_i)$ be the outwards neighborhood of the event E_i in the VDL graph, i.e.,

$$\Gamma(E_i) = \{E_j \mid E_j \text{ is VDL on } E_i\}.$$

The notion of VDL was inspired by the following *undirected* notion of Moser and Tardos [14]:

Definition 2 (Moser and Tardos [14])**.** *Let E_i, E_j be events, $i, j \in \{1, ..., m\}$. We say that E_i, E_j are lopsidependent if there exist two evaluations α, β, that differ only on variables in $e^i \cap e^j$, such that:*

1. *α makes E_i happen, β makes E_j happen and*
2. *either β makes $\overline{E_i}$ happen or α makes $\overline{E_j}$ happen.*

Notice again that an event E_i is never lopsidependent on itself, nor on any event whose scope shares no variables with e^i.

We will prove that two events are lopsidependent in the sense of Definition 2 if at least one is VDL on the other. Thus, the lopsidependency undirected graph of the events under the binary relation of Definition 2 is the *undirected* graph that underlies the VDL graph.

Claim 1. *Two events E_i, E_j, $i, j \in \{1, ..., m\}$, are lopsidependent (in the sense of definition 2), if and only if $E_j \in \Gamma(E_i)$ or $E_i \in \Gamma(E_j)$.*

Proof. ($\Rightarrow$) By definition 2, there exist evaluations α, β that differ only on variables in $e^i \cap e^j$ such that E_i occurs under α, E_j occurs under β and either $\overline{E_j}$ occurs under α or $\overline{E_i}$ occurs under β.

It is immediate to see that if $\overline{E_j}$ occurs under α, $E_j \in \Gamma(E_i)$ and that, if $\overline{E_i}$ occurs under β, $E_i \in \Gamma(E_j)$.

($\Leftarrow$) Let $E_j \in \Gamma(E_i)$. Then, there are two evaluations $\alpha = (a_1, ..., a_l), \beta = (b_1, ..., b_l)$ that differ only in e^i, such that E_i, $\overline{E_j}$ occur under α and E_j occurs under β. If evaluations α, β differed only in $e^i \cap e^j$, there would be nothing to prove.

Define evaluation $\beta' = (b'_1, ..., b'_l)$ such that:

- $b'_r = a_r$, for $r = 1, \ldots, l$ such that $X_r \notin e^i \cap e^j$ and
- $b'_r = b_r$, for $r = 1, \ldots, l$ such that $X_r \in e^i \cap e^j$.

Since α differs from β only on variables in e^i, it follows that β' differs from β only on variables in $e^i \setminus e^j$. Now, since E_j holds for β and does not depend on variables not in its scope, E_j holds under β' also. Thus, evaluations α, β' fulfill the requirements of Definition 2.

Similarly if $E_i \in \Gamma(E_j)$. □

We now give the (non-directed) classical definition of lopsidependency formulated by Erdős and Spencer [10]:

Definition 3 (Erdős and Spencer [10]). *Let $E_1, \ldots, E_m$ be events in an arbitrary probability space, G a graph on the indices $1, \ldots, m$. We say that G is an* Erdős-Spencer lopsidependency graph *(for the events) if*

$$\Pr\left[E_j \mid \bigcap_{i \in I} \overline{E_i}\right] \leq \Pr\left[E_j\right]$$

for all j, I with $j \notin I$ and no $i \in I$ adjacent to j.

Observe that if two events are (strictly) negatively correlated, they should be connected by an edge in the Erdős-Spencer lopsidependency graph. Since the notion of negative correlation is symmetric, we conclude that it would not make sense to define a directed notion of lopsidependency in the general framework without variables.

We will now show that the underlying undirected graph of the VDL graph (in other words the lopsided undirected graph in the sense of Definition 2) is an Erdős-Spencer dependency graph. This fact shows that it might be advantageous for applications of LLLL to use the notion of VDL rather than the classical undirected Erdős-Spencer lopsidependency or the undirected lopsidependency of Moser and Tardos [14] (see Theorem 1 below).

Lemma 1. *For any event E_j, $j = 1, ..., m$, let I be the set of indices of events not in $\Gamma(E_j) \cup \{E_j\}$. Then, it holds that:*

$$Pr\left[E_j \mid \bigcap_{i \in I} \overline{E_i}\right] \leq Pr[E_j].$$

Proof. Let $E_k = \bigcap_{i \in I} \overline{E_i}$. Since $E_i \notin \Gamma(E_j)$, for all $i \in I$, it holds that for any evaluation α that makes E_j, E_k hold, there is no evaluation β that differs from α only in e^j that makes $\overline{E_k}$ hold.

Now, for the events E_j, E_k, it holds that:

$$Pr[E_j \mid E_k] \leq Pr[E_j] \Leftrightarrow \frac{Pr[E_j \cap E_k]}{Pr[E_k]} \leq Pr[E_j] \Leftrightarrow$$
$$Pr[E_j \cap E_k] \leq Pr[E_j]Pr[E_k]$$

Suppose now $\alpha = (a_1, ..., a_l), \beta = (b_1, ..., b_l)$ are two evaluations obtained by independently sampling the random variables twice, once to get α and once to get β. It holds that:

$$Pr[\underbrace{(\alpha, \beta) \text{ are such that } \alpha \text{ makes } E_j, E_k \text{ happen}}_{\text{event } S}] = Pr[E_j \cap E_k].$$

Indeed, event S above imposes no restriction on β.

Let now $\alpha' = (a'_1, ..., a'_l)$, $\beta' = (b'_1, ..., b'_l)$ be two evaluations obtained by α, β by swapping values in variables in $e^k \setminus e^j$:

- $a'_i = b_i$, for all i such that $X_i \in e^k \setminus e^j$, $a'_i = a_i$ for the rest,
- $b'_i = a_i$, for all i such that $X_i \in e^k \setminus e^j$ and $b'_i = b_i$ for the rest.

Since $a'_i = a_i$ for all i such that $X_i \in e^j$, a' makes E_j happen. Observe also that the coordinates of β' that correspond to variables in e^k and whose values are different than those of α, are all in $e^j \cap e^k \subseteq e^j$. By the hypothesis, and since variables not in e^k do not influence E_k, if E_j, E_k hold under α, then E_k holds under β'.

Obviously α', β' are two independent samplings of all variables, since all individual variables were originally sampled independently, and we only changed the positioning of the individual variables. Thus, it holds that:

$$Pr[\underbrace{\alpha' \text{ makes } E_j \text{ happen and } \beta' \text{ makes } E_k \text{ happen}}_{\text{event } T}] = Pr[E_j]Pr[E_k].$$

Now, by the hypothesis and the construction of α', β', it also holds that S implies T. Thus:

$$Pr[S] \leq Pr[T] \Leftrightarrow Pr[E_j \cap E_k] \leq Pr[E_j]Pr[E_k]. \quad \square$$

Suppose now that the size of the outwards neighborhood of any event in the VDL graph is bounded by $d \in \mathbb{N}$, i.e. $max\{|\Gamma(E_j)| \mid 1 \leq j \leq m\} \leq d$ and that $p \in [0,1]$ is the maximum of the probabilities $Pr[E_j]$, $j = 1, ..., m$. In the next section we will prove (e is the base of the natural logarithms):

Theorem 1 (Symmetric Lopsided Local Lemma for VDL)**.** *If $ep(d+1) \leq 1$, then $Pr[\overline{E_1} \wedge \overline{E_2} \wedge \cdots \wedge \overline{E_m}] > 0$, i.e. there exists an evaluation of the variables X_i for which none of the events E_j hold.*

3 Algorithmic Lopsided Lovász Local Lemma

The notion of lopsidependency we use hereafter is the directed notion of Definition 1, which we called VDL. The outwards neighborhood operation of Γ is always in terms of VDL.

To prove Theorem 1 in our framework, we give a Moser-like algorithm, M-Algorithm below, and find an upper bound to the probability that it lasts for n phases.

It is trivial to see that if M-Algorithm ever stops, it returns an evaluation for which none of the "undesirable" events $E_1, ..., E_m$ holds.

We first prove:

Lemma 2. *Consider an arbitrary call of Resample(E_j). Let $\mathcal{X}$ be the set of events that do not occur at the start of this call. Then, if and when this call terminates, all events in $\mathcal{X} \cup \{E_j\}$ do not occur.*

Proof. Consider first an event E_s in $\mathcal{X}$ and assume Resample(E_j) terminates and that E_s occurs during the execution of Resmaple(E_j). Thus, E_s occurs at step 1 of a Resample call during the execution of Resample(E_j). Suppose that Resample(E_k) is the last call of a Resample made during Resample(E_j) when E_s turns from non-occurring to occurring and stays occurring until Resample(E_j) has ended (E_k could be E_j itself).

Consider the evaluation just before Resample(E_k) was called, call it α, and the one just after Resample(E_k) was called, call it β. Under α, E_k occurred (else Resample(E_k) would not have been called at this point), and also E_s did

M-Algorithm

1. Sample the variables X_i, $i = 1, ..., l$ and let α be the resulting evaluation.
2. **while** there exists an event that occurs under the current evaluation, let E_j be the least indexed such event and **do**
3. Resample(E_j)
4. **end while**
5. Output current assignment α.

Resample(E_j)

1. Resample the variables in e^j.
2. **while** some event in $\Gamma(E_j) \cup \{E_j\}$ occurs under the current evaluation, let E_k be the least indexed such event and **do**
3. Resample(E_k)
4. **end while**

not occur (by hypothesis). Now, since β is the evaluation right after the call of Resample(E_k), the evaluations α and β differ only on the variables in e^k and, since β makes E_s occur, we have that $E_s \in \Gamma(E_k) \cup \{E_k\}$. This means that Resample(E_k) would not terminate as long as E_s occurred. Contradiction.

Secondly, the proof that E_j does not occur at the end of Resample(E_j) is immediate. □

A *root call* of Resample is any call of Resample made when executing line 3 of M-Algorithm. On the other hand, a *recursive call* of Resample is a call made from within another call of Resample. A *phase* is the execution period from the start until the end of a Resample root call.

By Lemma 2, we know that the events of the root calls of Resample are pairwise distinct. Therefore:

Corollary 1. *There are at most m phases in any execution of M-Algorithm.*

Consider now *rooted forests*, i.e. forests of trees such that each tree has a special node designated as its root.

The nodes of rooted forests are labeled by the events E_j, $j = 1, ..., m$, with repetitions of the labels allowed and they are ordered as follows: children of the same node are ordered as their labels are; nodes in the same tree are ordered by preorder (respecting the ordering between siblings) and finally if the label on the root of a tree T_1 precedes the label of the root of T_2, all nodes of T_1 precede all nodes of T_2.

The number of nodes of a forest $\mathcal{F}$ is denoted by $|\mathcal{F}|$.

Definition 4. *A labeled rooted forest $\mathcal{F}$ is called* feasible *if the following conditions hold:*

1. *The labels of the roots of $\mathcal{F}$ are pairwise distinct.*
2. *If u, v are siblings (have common parent), then the labels of u, v are distinct.*
3. *Let E_i, E_j be the labels of nodes u, v respectively, where u is a child of v. Then, $E_i \in \Gamma(E_j) \cup \{E_j\}$.*

Now consider an execution of M-Algorithm that lasts for n phases. We construct in a unique way, depending on the steps, a feasible forest with n nodes as follows:

1. The forest under construction will have as many roots as root calls of Resample. These roots will be labeled by the event of the corresponding root call.
2. A tree that corresponds to a root call Resample(E_j) will have as many non-root nodes as the number of recursive calls of Resample within Resample(E_j). The non-root nodes will be labeled by the events of those recursive calls.
3. The non-root nodes are organized within the tree with root-label E_j so that a node that corresponds to a call Resample(E_k) is parent to a root that corresponds to a call Resample(E_l), if Resample(E_l) appears immediately on top of Resample(E_k) in the recursive stack that implements the root call Resample(E_j).

It is straightforward to check, inspecting the succession of steps of M-Algorithm in an execution, and making use of Lemma 2, that a forest constructed as above from the consecutive phases of the execution of M-Algorithm is indeed a feasible forest in the sense of Definition 4. It is not true however that every feasible forest corresponds to the consecutive phases of some execution of M-Algorithm. For example a feasible forest where the child of a node v has a label that differs from the label of v and is not lopsidependent on it (although could be dependent on it) can never be constructed as above.

Definition 5. *The* witness forest *of an execution of the M-Algorithm is the feasible forest constructed as described above. Given a feasible forest $\mathcal{F}$ with n nodes, we denote by $W_{\mathcal{F}}$ the event that M-Algorithm lasts for n phases and during these n phases, $\mathcal{F}$ is constructed.*

We also set:

$$\hat{P}_n := \Pr\left[\bigcup_{\mathcal{F}:|\mathcal{F}|=n} W_{\mathcal{F}}\right] = \sum_{\mathcal{F}:|\mathcal{F}|=n} \Pr\left[W_{\mathcal{F}}\right], \tag{1}$$

where the last equality holds, because the events $W_{\mathcal{F}}$ are disjoint. Obviously now:

$$Pr[\text{M-Algorithm lasts for } n \text{ phases}] = \hat{P}_n. \tag{2}$$

ValAlg (**Input:** feasible forest $\mathcal{F}$ with labels $E_{j_1}, ..., E_{j_n}$)

1. Sample the variables X_i, $i = 1, ..., l$.
2. **for** $s = 1, ..., n$ **do**
3. **if** E_{j_s} does not occur under current evaluation **then**
4. **return** failure and exit
5. **else**
6. Resample the variables in e^{j_s}.
7. **end if**
8. **end for**
9. **return** success.

To find an upper bound for $\hat{P}_n$, consider the following algorithm, which we call *Validation Algorithm*:

The validation algorithm ValAlg, takes as input a feasible forest $\mathcal{F}$ with n nodes, labeled with the events $E_{j_1}, ..., E_{j_n}$ (ordered as their respective nodes) and outputs a Boolean value *success* or *failure*.

Intuitively, with input a feasible forest with labels $E_{j_1}, ..., E_{j_n}$, ValAlg initially generates a random sampling of the variables, then checks the current event and, if it holds, resamples its variables and goes to the next event. If the algorithm manages to go through all events, it returns *success*, otherwise, at the first event that does not hold under the current assignment, it returns *failure* and stops.

We first give a result about the probability distribution of the assignment to the variables X_i during the execution of ValAlg.

Lemma 3 (Randomness Lemma)**.** *At the beginning of any iteration of the* ***for*** *loop of Step* 2 *of ValAlg, the distribution of the current assignment of values to the variables X_i, $i = 1, ..., l$ is as if all variables have been sampled anew. Therefore the probability of any event occurring at such an instant is bounded from above by p.*

Proof. The result follows from the fact that if at the previous iteration of the **for** loop of Step 2 the event E_{j_s} has been checked, only the values of the variables in e^{j_s} have been exposed, and these are resampled anew. □

Definition 6. *Given a feasible forest $\mathcal{F}$ with n nodes, we say that $\mathcal{F}$ is validated by ValAlg if the latter returns* success *on input $\mathcal{F}$. The event of this happening is denoted by V_F. We also set:*

$$P_n = \sum_{\mathcal{F}:\ |\mathcal{F}|=n} \Pr[V_{\mathcal{F}}]. \tag{3}$$

We now claim:

Lemma 4. *For any feasible forest $\mathcal{F}$, the event $W_{\mathcal{F}}$ implies the event $V_{\mathcal{F}}$, therefore:*

$$\hat{P}_n \leq P_n \tag{4}$$

Proof. Indeed, if the random choices made by an execution of M-Algorithm that produces as witness forest $\mathcal{F}$ are made by ValAlg on input $\mathcal{F}$, then clearly ValAlgo will return "success". □

Lemma 5. *For any feasible forest $\mathcal{F}$ with n nodes, $\Pr[V_{\mathcal{F}}]$ is at most p^n.*

Proof. Immediate corollary of the Randomness Lemma 3. □

Therefore to find an upper bound for $P_n = \sum_{\mathcal{F}:\ |\mathcal{F}|=n} \Pr[V_{\mathcal{F}}]$, it suffices to find an upper bound on the number of feasible forests with n nodes, a fairly easy exercise in enumerative combinatorics, whose solution we outline below.

First to any feasible forest $\mathcal{F}$ with n nodes we add new leaves to get a new labeled forest $\mathcal{F}'$ (perhaps not feasible any more) comprising of m full $(d+1)$-ary trees whose internal nodes comprise the set of all nodes (internal or not) of $\mathcal{F}$ (recall m is the total number of events $E_1, \ldots, E_m$). Specifically, we perform all the following additions of leaves:

1. Add to $\mathcal{F}$ new trees, each consisting of a single root/leaf labeled with a suitable event, so that the set of labels of all roots of $\mathcal{F}'$ is equal to the the set of all events $E_1, \ldots, E_m$. In other words, the labels of the added roots/leaves are the events missing from the list of labels of the roots of $\mathcal{F}$.

2. Hang from all nodes (internal or not) of $\mathcal{F}$ new leaves, labeled with suitable events, so that the set of the labels of the children in $\mathcal{F}'$ of a node u of $\mathcal{F}$ labeled with E is equal to the set of the events in $\Gamma(E) \cup \{E\}$. In other words, the labels of the new leaves hanging from u are the events in $\Gamma(E) \cup \{E\}$ that were missing for the labels of the children of u in $\mathcal{F}$.

3. Add further leaves to all nodes u of $\mathcal{F}$ so that every node of $\mathcal{F}$ (internal or not in $\mathcal{F}$) has exactly $d+1$ children in $\mathcal{F}'$; label all these leaves with the first $(d+1) - |\Gamma(E) \cup \{E\}|$ events from the list $E_1, \ldots, E_m$, so that labels of siblings are distinct, and therefore siblings can be ordered by the index of their labels.

Notice first that indeed, the internal nodes of $\mathcal{F}'$ comprise the set of all nodes (internal or not) of $\mathcal{F}$.

Also, the labels of the nodes of a labeled forest $\mathcal{F}'$ obtained as above are uniquely determined from the rooted planar forest structure of $\mathcal{F}'$ when labels are ignored, but the ordering of the nodes imposed by them is retained (a forest comprised of rooted trees is called *rooted planar* if the roots are ordered and if the children of each internal node are ordered). Indeed, the roots of $\mathcal{F}'$ are labeled by $E_1, \ldots, E_m$; moreover, once the the label of an internal node of $\mathcal{F}'$ is given, the labels of its children are distinct and uniquely specified by steps (2) and (3) of the construction above.

Finally, distinct $\mathcal{F}$ give rise to distinct $\mathcal{F}'$.

By the above remarks, to find an upper bound to the number of feasible forests with n nodes (internal or not), it suffices to find an upper bound on the number of rooted planar forests with n internal nodes comprised of m full $(d+1)$-ary rooted planar trees.

It is well known that the number t_n of full $(d+1)$-ary rooted planar trees with n internal nodes is equal to $\frac{1}{dn+1}\binom{(d+1)n}{n}$, see e.g. [21, Theorem 5.13]. Now by Stirling's approximation easily follows that for some constant $A > 1$, depending only on d, we have:

$$t_n < A\left(\left(1+\frac{1}{d}\right)^d p(d+1)\right)^n. \tag{5}$$

Also obviously the number f_n of rooted planar forests with n internal nodes that are comprised of m $(d+1)$-ary rooted planar trees is given by:

$$f_n = \sum_{\substack{n_1+\cdots+n_m=n \\ n_1,\ldots,n_m\geq 0}} t_{n_1}\cdots t_{n_m}. \tag{6}$$

From (5) and (6) we get:

$$f_n < (An)^m\left(\left(1+\frac{1}{d}\right)^d p(d+1)\right)^n < (An)^m(ep(d+1))^n. \tag{7}$$

So by Lemma 4, equation (3), Lemma 5 and equation (6) we get the following Theorem, which is essentially a detailed restatement of Theorem 1:

Theorem 2. *Assuming p and d are constants such that $\left(1+\frac{1}{d}\right)^d p(d+1) < 1$, (and therefore if $ep(d+1) \leq 1$), there exists an integer N, which depends linearly on m, and a constant $c \in (0,1)$ (depending on p and d) such that if $n/\log n \geq N$ then the probability that M-Algorithm lasts for n calls of Resample is $< c^n$.*

Clearly, when M-algorithm stops we have found an evaluation such that none of the events occurs. Since, by the above Theorem, this happens with probability close to 1 for large enough n, Theorem 1 follows.

References

[1] Paul Erdős and László Lovász. Problems and results on 3-chromatic hypergraphs and some related questions. *Infinite and finite sets*, 10:609–627, 1975.

[2] M. Szegedy. The lovász local lemma - a survey. In Springer, editor, *Computer Science - Theory and Applications*, pages 11–23, 2013.

[3] Robin A. Moser. A constructive proof of the Lovász local lemma. In *Proceedings of the 41st annual ACM Symposium on Theory of Computing*, pages 343–350, 2009.

[4] Noga Alon. A parallel algorithmic version of the local lemma. *Random Structures & Algorithms*, 2(4):367–378, 1991.

[5] József Beck. An algorithmic approach to the Lovász local lemma. I. *Random Structures & Algorithms*, 2(4):343–365, 1991.

[6] Aravind Srivinsan. Improved algorithmic versions of the Lovász local lemma. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 611–620. Society for Industrial and Applied Mathematics, 2008.

[7] Terence Tao. Moser's entropy compression argument. 2009. Available: https://terrytao.wordpress.com/2009/08/05/mosers-entropy-compression-argument/.

[8] Joel Spencer. Robin Moser makes Lovász local lemma algorithmic!, 2010.

[9] Ioannis Giotis, Lefteris Kirousis, Kostas I. Psaromiligkos, and Dimitrios M. Thilikos. On the algorithmic Lovász local lemma and acyclic edge coloring. In *Proceedings of the twelfth workshop on analytic algorithmics and combinatorics*. Society for Industrial and Applied Mathematics, 2015. http://epubs.siam.org/doi/pdf/10.1137/1.9781611973761.2.

[10] Paul Erdős and Joel Spencer. Lopsided Lovász local lemma and Latin transversals. *Discrete Applied Mathematics*, 30(2-3):151–154, 1991.

[11] Paul R. Berman, AD Scott, and M Karpinski. Approximation hardness and satisfiability of bounded occurrence instances of SAT. *ECCC*, 10(022), 2003.

[12] Heidi Gebauer, Robin A. Moser, Dominik Scheder, and Emo Welzl. The Lovász local lemma and satisfiability. In *Efficient Algorithms*, pages 30–54. Springer, 2009.

[13] Heidi Gebauer, Tibor Szabó, and Gábor Tardos. The local lemma is tight for SAT. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 664–674. SIAM, 2011.

[14] Robin A. Moser and Gábor Tardos. A constructive proof of the general Lovász local lemma. *Journal of the ACM (JACM)*, 57(2):11, 2010.

[15] Kashyap Kolipaka, Babu Rao, and Mario Szegedy. Moser and Tardos meet Lovász. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 235–244. ACM, 2011.

[16] James B. Shearer. On a problem of Spencer. *Combinatorica*, 5(3):241–245, 1985.

[17] Nicholas Harvey and Jan Vondrák. An algorithmic proof of the lopsided Lovász local lemma. arXiv:1504.02044.

[18] David G Harris. Lopsidependency in the Moser-Tardos framework: beyond the lopsided Lovász local lemma. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1792–1808. SIAM, 2015.

[19] Joel Spencer. Ten lectures on the probabilistic method. In *SIAM*, volume 64, 1994.

[20] Dimitris Achlioptas and Fotis Iliopoulos. Random walks that find perfect objects and the Lovász local lemma. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 494–503. IEEE, 2014.

[21] Robert Sedgewick and Philippe Flajolet. *An introduction to the analysis of algorithms*. Addison-Wesley, 2013.